

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE		3. REPORT TYPE AND DATES COVERED
		15 November 2001		Final Technical Report, April 1998 – March 2001
4. TITLE AND SUBTITLE			5. FUNDING NUMBERS	
Development of Design Methodologies for Robust, Adaptive Autonomous Hierarchical Systems: Final Report			N00014-98-1-0535 02PR02522-00	
6. AUTHOR(S)				
B.D.O. Anderson, T.S. Brinsmead, J. Davoren, T. Moor, M. Smith, S. Su				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)				
Australian National University Department of Systems Engineering Acton Act 0200, Australia				
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
Office of Naval Research (ONR) Ballston Centre Tower One 800 North Quincy Street Arlington, VA 22217-56600				
11. SUPPLEMENTARY NOTES				
27 Pages. Name of responsible person – Thomas Brinsmead, telephone number +61-2-6125-8674				
12a. DISTRIBUTION/AVAILABILITY STATEMENT			12b. DISTRIBUTION CODE	
Approved for public release, distribution is unlimited.				
ABSTRACT (Maximum 200 words)				
The report describes the research outcomes from a three year study of methodologies to secure robust behaviour of nonlinear systems with high speed adaptation capability and principles for control architectures with self-organizing adaptive capabilities appropriate for use in intelligent, autonomous, unmanned air vehicles.				
14. SUBJECT TERMS			15. NUMBER OF PAGES	
ONR, Australia, Unmanned air vehicles, Hierarchical control architecture, Robust, Adaptive control methodologies				
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL	

20020110 045

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 15-11-2001		2. REPORT DATE FINAL TECHNICAL REPORT		3. DATES COVERED (From - To) APR 1998 - MAR 2001	
4. TITLE AND SUBTITLE DEVELOPMENT OF DESIGN METHODOLOGIES FOR ROBUST ADAPTIVE AUTONOMOUS HIERARCHICAL SYSTEMS : FINAL REPORT				5a. CONTRACT NUMBER N00014-98-1-0535	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
				5d. PROJECT NUMBER 02PR02522-00	
6. AUTHOR(S) B D O ANDERSON, T S BRINSMEAD, J DAVOREN, T MOOR, M SMITH, S SU				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AUSTRALIAN NATIONAL UNIVERSITY, DEPARTMENT OF SYSTEMS ENGINEERING ACTON ACT 0200 AUSTRALIA				8. PERFORMING ORGANIZATION REPORT NUMBER NA	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) OFFICE OF NAVAL RESEARCH BALLSTON CENTRE TOWER ONE 800 NORTH QUINCY STREET ARLINGTON VA 22217-5660				10. SPONSOR/MONITOR'S ACRONYM(S) ONR	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER	
12. DISTRIBUTION AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE, DISTRIBUTION IS UNLIMITED					
13. SUPPLEMENTARY NOTES NA					
14. ABSTRACT The report describes the research outcomes from a three year study of methodologies to secure robust behaviour of nonlinear systems with high speed adaptation capability, and principles for control architectures with self-organising adaptive capabilities appropriate for use in intelligent autonomous, unmanned air vehicles.					
15. SUBJECT TERMS Unmanned Air Vehicles, Hierarchical Control Architecture, Robust and Adaptive Control Methodologies.					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 27	19a. NAME OF RESPONSIBLE PERSON THOMAS BRINSMEAD
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code) +61-2-6125 8674

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g., 30-06-1998; xx-08-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATES COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33615-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report, e.g. 1F665702D1257.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. AFOSR-82-1234.

5d. PROJECT NUMBER. Enter all project numbers as they appear in the report, e.g. 1F665702D1257; ILIR.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORS AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

Development of Design Methodologies for Robust, Adaptive Autonomous Hierarchical Systems: Final Report

B. D. O Anderson, T. S. Brinsmead, J. Davoren, T. Moor, M. Smith, S.Su

15 November 2001

1 Introduction

ONR contract number N00014-98-1-0535 commenced on 1 April 1998, funding research performed by the Australian National University, Department of Systems Engineering. This final report summarizes the research outcomes of the contract, including research previously reported in Interim Progress Reports [PR01] and [PR02].

The technical framework of the contract centred on methodologies to secure robust behaviour of nonlinear systems with high speed adaptation capability, and principles for control architectures with self-organising adaptive capabilities appropriate for use in intelligent autonomous unmanned air vehicles.

A research program was set up to focus on two key topics:

- Low level behaviours: Development of identification and control methods to deal with time-varying and/or highly nonlinear uncertain systems, while providing adaptivity and robustness.

In particular, the main operational requirement for the lower levels of a controlled hierarchy is that each component be capable of responding to a command issued by high level components in an effective way. At the same time, however, the lower level components need to be autonomous, given that they are operating in an uncertain and rapidly changing environment. Consequently, research under this contract focussed on identification and control methods that provide adaptivity with robustness.

- High level behaviours: Development of control methodologies for hierarchical systems.

In particular, a hierarchically organized system comprises a collection of sub-systems. Sensor information flows upwards and control decisions and command strategies flow downwards to be implemented locally. Fundamental rules and principles are needed to understand such issues as how many layers of the hierarchy are optimal. Should the time scale of operation or some other criterion govern the level at which a sub-system is located in a hierarchy? How should information overload on upper hierarchical levels be avoided, that is, how should the systems communicate and what type of information should be shared in other words, how should the various levels of the hierarchy be coupled? Which decisions can and/or should be centralized rather than decentralized? How should adaptive capabilities be distributed through the hierarchy?

AQ F02-04-0479

This final report summarizes the research outcomes of the contract for the three years from 1998/9-2000/01. It is organized as follows. Section 2 revisits the relevance of the work to the Navy and summarizes some of the scientific challenges that had to be overcome during the course of the contract. In Section 3 we review the research objectives stated in the original research proposal and describe how these objectives were updated during the course of the contract. In the following sections we then summarise in detail the research contract outcomes which comprised two main components low level and high level behaviour. Publications that have been produced in the course of the contract are referenced as [XXXnn], where XXX is a two or three character alpha code corresponding to the research area and nn is numeric. Citations of publications that were not produced as part of research conducted under this contract are in the format [Author (year)]. A reference list appears as a final section after the conclusion and list of research team members.

2 Background

2.1 Rationale for the Navy

There are very few design methodologies for control systems that integrate elements of robustness, adaptivity, autonomy and hierarchical control even for linear systems. Nonlinear methodologies are even less advanced, although there are some exceptions [Krstic et al. (1995)], [Jiang and Praly (1998)]. An unmanned combat air vehicle is certainly an example of a system that requires such functionality. Specifically, the Navy wants to use vehicles which:

- can operate after damage and in all weathers, which gives a requirement for robustness;
- might have to change control strategy in response to environmental changes such as battle damage, release of stores, or single/multiple vehicle operation, which gives rise to a requirement for adaptivity;
- are unmanned, and have greater manoeuvre capability than straight and level flight on a pre-planned route, perhaps having the ability to evade a hostile threat, or to acquire and attack a target, which gives rise to requirement for autonomy;
- can operate either in formation or separately, which gives rise to a requirement for hierarchical and/or cooperative control.

2.2 Scientific Roadblocks

Most research contracts can be expected to identify and overcome scientific roadblocks in order to secure a customer-specified end objective. This contract is no exception. In this section, we describe some of these scientific roadblocks.

The great bulk of control systems design methodologies are directed at linear systems (for example, see [Morari and Zafiriou (1989)], [Zhou et al. (1996)], [Franklin et al. (1994)], [Kuo (1991)]). Methodologies are embodied both in algorithms, as well as in a hugely rich conceptual framework that allows designers to confidently use algorithms to design controllers that are eventually implemented in real systems, despite the fact that exact modelling of real systems is impossible. In broad terms, it is fair to say that the situation for nonlinear systems is nowhere near as advanced as

it is for linear systems (although advances in nonlinear systems have been especially great over the last decade or two). In particular:

- the vast majority of control methodologies and algorithms, but of course not all, are for linear systems;
- there is much less known about the robust control of nonlinear systems than that for linear systems;
- the non-specifically-linear-system conceptual content of adaptive control ideas is just being isolated. It is likely that this will provide a basis for successful nonlinear adaptive control. Of course, some nonlinear adaptive control techniques, principally embodied in algorithms, are available. However in general, such algorithms have less associated conceptual content than those for linear systems;
- even for linear systems, the boundary between the situations in which it is appropriate to implement robust rather than adaptive control, is indistinct;
- nearly all control methodology is focused on low level (although often critical) tasks, such as pitch control for an aircraft. In contrast, little is known about hierarchical or cooperative control. This is relevant to, for example, integrating auto-pilots, guidance algorithms and threat evaluation, including possibly costly measurement strategies, such as, for example, switching on of a radar, or flight-path deviation in order to improve visual targeting;
- conventional control theories for the design of hierarchical systems are primitive at best. Discrete event systems concepts are applicable to some problems, but discrete event modelling of physical systems is not at all straightforward.

3. Adjustment of the Objectives in the Course of the Contract

The original research proposed concentrated on two broad areas, low level behaviour which focussed on adaptive nonlinear control methodology and high level behaviour. The originally proposed research for low level behaviour included components of iterative controller optimization for nonlinear systems; and the problem of re-configurability. High level behaviours were to include hierarchical control—the importance of time scale and the duality of control and information; and sensor fusion.

During the course of the contract there was a shift in emphasis in research on the low level behaviour, motivated partly by suggestion of the ONR contract monitor towards multiple model adaptive control and associated issues. Research on high level behaviours was redirected towards Discrete Event Systems and Hybrid Systems in order to take advantage of some fortuitous research contacts including Professor S. Sastry from UC Berkeley. Consequently, work was redirected to bring methods of discrete event and hybrid systems theory to bear on hierarchical control problems. Work in this area was later extended to also include development of logics for hybrid systems that allow for robustness analysis.

As well as technical drivers for variation in the research contract, there were also some financial issues that impacted on the acquisition of research personnel. These will be described in more detail in the following sections.

3.1 Technical Drivers

During the course of the contract, the chief investigator, as well as various associates attended program review meetings, and undertook discussions with Dr Allen Moshfegh. In the course of these discussions, Dr Moshfegh requested that some alternative approaches be considered.

In the area of low level behaviour, for example, attention was turned to multiple model adaptive control, a decision that was justified on several grounds, including the suggestion by Dr Moshfegh of its intrinsic interest to the Navy. Other reasons were

- the relative ease with which multiple model adaptive algorithms rather than conventional algorithms based on parameter estimation can be extended from the linear to the nonlinear case;
- an objective specified in the original contract of handling the problem of re-configurability.

In the area of high level behaviours, attention was turned to the possible application of hybrid systems theory, and the use of logic and other formal methods to design hybrid systems controllers. A crucial feature of these logics is the ability to rigorously specify and analyse robustness properties for hybrid systems. Furthermore, in order to develop methodology for solving filtering and smoothing problems associated with discrete-state systems, more research than had originally been planned was carried out on hidden Markov models. This change in research direction occurred because both multiple model adaptive controllers and hybrid systems are typically characterized by discrete, as well as continuous states.

3.2 Financial Drivers of changes in research direction

The originally budget called for approximately 200K to be paid for the three years from 1998/99-2000/01. However, this sum was paid for only the first year, 1998/99, with approximately \$150K provided in 1999/2000. A still lesser sum was provided in total for 2000/01. A final instalment of approximately \$20K is currently expected. The reduced funding has impacted on the work on high-level behaviour to a greater degree than the work on low level behaviours.

3.3 Research Outcomes: Detail

In the following two sections, we describe the contract research outcomes in more detail. Section 4 reports the outcomes of research on low level behaviours, and the Section 5 reports on high level behaviour.

Publications that were produced during the course of the project include some fourteen journal articles, twenty-five conference papers, four book chapters and six technical reports. Some articles have been accepted and are pending publication, and the technical reports have been submitted for publication and are undergoing review. In addition, various fragments of computer code were generated to emphasise the principal applicability of our results. Some of the attached publications illustrate theoretical insights by simulations, and, in the scope of selected examples, some of the controller designs involve some computer code: see, for example [HL02, HL04, HL06, HL07]. However, no ready-to-go software modules were promised nor have been developed in the context of this research contract.

4. Low Level Behaviour

The original contract broke the low level work into three areas: adaptive nonlinear control methodology; iterative controller optimization for nonlinear systems; and the problem of re-configurability. Subsequent sections report against slightly different headings, given the evolution of the work, and the redirections made. The subsequent headings are:

- Nonlinear closed-loop identification;
- Controller design for nonlinear systems;
- Safe adaptive control, and iterative feedback tuning;
- Multiple model adaptive control;
- Hidden Markov Models.

These new categories are closely related to the old ones. In particular, the problem of adaptive nonlinear control methodology was split into two sub-problems: the identification of nonlinear systems and controller design for nonlinear systems. The research on reconfigurability was split into research on multiple model adaptive control and on hidden Markov Models.

We now describe the research on nonlinear identification. Many adaptive control algorithms require there to be, embedded implicitly within the algorithms, an identification of the unknown plant. This identification is then used, often on a certainty equivalent basis, to design a controller. The controller is changed as the identifier updates the plant model. The task of identification and controller design is treated in the following two sections.

The first of these, Section 4.2, deals directly with nonlinear design methodologies. Here, our aim is to expand the toolbox available for the design of nonlinear systems, noting that this expansion may well be within the context of nonlinear adaptive control.

Section 4.3 deals with iterative controller optimization (iterative feedback tuning), while exploring fundamental safety issues in adaptive control which were originally investigated for linear systems. Although linearity of the system is mostly irrelevant, these issues needed to be understood in order to put the nonlinear adaptive control methodologies in a logical framework. Some of these issues were able to be immediately extended to nonlinear systems, while for others, the nonlinear domain had to be reached via thorough examination of the ideas for linear systems. About half of the research reported in this section however, is directly nonlinear.

Multiple model adaptive control is dealt with in the Section 4.4 and the following section deals with hidden Markov models, principally for reasons indicated earlier. The most important of these

reasons is that hidden Markov Models can very conveniently model the behaviour of discrete state systems. High level behaviours have been treated by hybrid systems methods, which involve both discrete and continuous states. In addition, multiple model adaptive control also involves both discrete and continuous states. Hidden Markov models provide one of the analysis tools for investigating hierarchical systems such as those that arise in multiple model adaptive control, potentially allowing filtering, prediction and smoothing.

4.1 Nonlinear closed-loop identification

The development of algorithms for plant identification, [Ljung(1987)], [Sjöberg et al.(1995)] in closed loop, even in the linear case, has been an important line of research occupying the attention of many people in the last few years [Ljung and Forsell (1999)], [Forsell and Ljung (1999)], [Van den Hof and Schrama (1995)]. The research has been motivated by several factors. Some of these are as follows.

- In a number of situations, identification in an open loop is difficult, or is simply not feasible. This occurs, for example, when the plant is unstable in open loop operation, including when it has an integrator or has significant open loop drift.
- There may be a controller already in the loop, which is to be re-tuned after improved plant identification.
- Closed-loop rather than open loop identification offers the possibility of capturing dynamic characteristics of the plant model that are critical for (closed-loop) control design.

Recent advances in closed-loop identification in the linear case and the fact that adaptive control algorithms frequently include identification as one component of the algorithm suggested the need to look at nonlinear closed-loop identification algorithms from the start of the contract.

Before describing the outcomes of the work, we explain why closed-loop identification is harder than open-loop identification: see also again [Van den Hof and Schrama (1995)]. There are at least two difficulties.

- The plant input signal and the disturbance are correlated, due to the feedback via the controller. If open-loop identification techniques are used, this correlation may bias the estimate.
- Even when the plant and controller are themselves linear, the closed-loop operator is a nonlinear function of the plant. Thus, even in the linear case, even once a closed operator is identified, and with knowledge of the controller, inverting this nonlinear relationship in order to determine the plant model can be difficult.

Of course, both these difficulties arise whether the plant is linear or nonlinear.

Earlier research by various workers, including ourselves, has overcome such difficulties for linear models. The abundance of nonlinearities in Navy related problems strongly motivates the need to extend standard linear theories to deal with the nonlinear issues. To this point, however, closed-loop identification of nonlinear systems has been left relatively untouched in the general literature. We have given it intensive effort within this contract. Note however, that at the end of this section, we indicate two further related problems for linear systems in which research was undertaken as part of the contract.

We have used two very broad approaches in nonlinear closed-loop identification. The first extends a linear systems identification algorithm, known in the "identification for control" [Gevers et al. (1999)] community as the Hansen scheme [Hansen et al. (1989)], and which uses coprime factors, to nonlinear systems.

We describe a coprime factor modelling approach [Vidyasagar (1985)] as follows. The transfer function of a time-invariant linear system may be expressed as the ratio of coprime polynomials. Alternatively, a rational transfer function can be expressed as a ratio of two coprime stable transfer functions, for example $1/(s-1) = [1/(s+\alpha)] [(s-1)/(s+\alpha)]^{-1}$. Many nonlinear operators can also be expressed as a cascade of operators, for example, ND^{-1} or $(D^{-})^{-1} N^{-}$ where D, N, D^{-}, N^{-} are stable operators, satisfying again a technical coprimeness condition. In the linear case, the Hansen scheme relies on the ability to parameterize an unknown plant in terms of the stable coprime factors of a known nominal model and the known controller, along with an unknown so-called Youla-Kucera parameter [Youla et al. (1976a)], [Youla et al. (1976b)] associated with the plant. The controller is assumed to stabilize both the true plant and the nominal model.

Rather than identifying the plant one identifies the Youla-Kucera parameter. The main advantage of this method is that crucially and non-obviously, this results in a closed-loop identification problem being transformed into one which is open-loop in nature. Our work has extended this method to work with nonlinear plants, a nonlinear nominal model and a nonlinear controller, that is, we have achieved the largest level of generality. See [ID01] for a survey introduction. Within this research, a series of technical issues have had to be addressed.

- A special type of coprimeness, termed differential coprimeness, was developed in order to account for the fact that nonlinear operators do not possess the distributivity property $A(B+C) = AB+AC$. See [ID01]-[ID05], which is related to earlier work by the principal investigator [Dasgupta and Anderson (1996)].
- While the definition of a right coprime factor representation, that is one of the form ND^{-1} , is relatively easily extended from the linear to the nonlinear case, this is not so for a left coprime realization, that is, one of the form $D^{-1}N$. In this case, a so-called kernel representation must be used instead. In the linear case, the kernel representation reduces to a left coprime realization. The papers [ID01]-[ID05] use kernel representations.
- Coprimeness requires both that the closed-loop input-output operator be bounded, and that the representation of systems by fractions involve bounded operators, that is, bounded inputs produce bounded outputs. Differential coprimeness requires that, in addition, the operators are continuous, that is, that small changes in inputs produce small changes in outputs. Assumption of such a continuity property is very reasonable for many engineering systems, although it cannot always be guaranteed.

We also investigated a second nonlinear closed-loop identification method embraced by the rather generic name "tailor-made approach". By exploiting knowledge of the controller, it minimizes the error between the measured closed-loop output of the true system, and the closed-loop output of the loop comprising a model with an adjustable parameter. At each instant of time, the adjustable parameter is set to the best current estimate, and a gradient scheme is used to update the parameter in order to reduce the closed-loop error.

A number of technical problems were encountered in the nonlinear implementation of such algorithms. The greatest of these involved the generation of the necessary gradients. Although in some cases, formal expressions for these gradients are available, they can involve unstable operators, and so are not usable in practice. The secondary difficulty is the requirement that both the

plant and the closed-loop performance must depend continuously on the parameters, in order that the gradient be well defined. A third restriction is that there can be collapse in the performance in a low signal to noise ratio environment. Such performance collapse in the presence of noise has been observed in many nonlinear systems algorithms, such as the phase locked-loop. This is a fundamental limitation for any algorithm using noisy gradients, such as the Extended Kalman Filter [Anderson and Moore (1979)] . Work on this material is outlined in [ID05], which presents a recursive identification method for a nonlinear plant operating in closed-loop with a nonlinear controller. The outputs of the plant are not a linear function of the unknown parameters in the situation considered, in contrast to a great many linear system parameter identification problems, and this is a superficial complicating feature.

Some nonlinear identifiers may use a batch or off-line procedure, whereas other identifiers continuously update the parameter estimate. In general, iterative (batch) algorithms present fewer technical difficulties than recursive ones which require additional stability issues to be resolved. Most of our work has focused on recursive algorithms. We have extended stability analysis of recursive identification algorithms for linear systems to the corresponding algorithms for nonlinear systems by employing passivity concepts [Sepulchre et al. (1996)], [Van der Schaft (1996)]. Passivity, fortunately, is not inherently a linear concept. Although in many respects it is difficult to find passive nonlinear operators, note that a nonlinear operator of the form $I + K$ where K has an induced norm less than unity, will necessarily be passive. See references [ID09]-[ID12] for research on recursive nonlinear identification. In contrast, reference [ID13] treats off-line or batch identification.

Lastly, in this subsection we describe results that deal with awkward issues that arise in the identification of linear systems which are equally relevant to nonlinear systems. Closed-loop identification in the presence of an unstable or a non-minimum phase controller is analysed in [ID14]: both these circumstances are possible in the nonlinear case. This paper shows that special precautions are needed when applying closed-loop identification methods in such circumstances, particularly if the identified model will be used to design a new controller.

Another paper, [ID15] deals with combining features of Hansen closed-loop identification schemes and tailor-made identification schemes. A major drawback of the Hansen scheme is that the order of the resulting model is not able to be tuned easily. A procedure which in a sense lies between Hansen and tailor-made is described in [ID15]. This hybrid procedure has the advantage of allowing the order of the resulting model to be tuned. Extending this to the nonlinear case is likely to be challenging, but potentially quite important.

4.2 Nonlinear control design methodology

During the contract we reviewed a number of nonlinear control methods, in the expectation that many nonlinear adaptive control algorithms would combine identification to obtain a plant model, with a controller design procedure to obtain a controller that is suitable for both the identified model and the actual plant. We completed several works in this area.

The first of these advocated the use of integrators to suppress constant disturbances and to track a constant reference input with zero error. Even in the linear case, the H-infinity control problem to secure these objectives is nonstandard. The same is true in the nonlinear case. In the linear case, however, special devices allow linear H-infinity theory to be applied to this non-standard problem [Mita et al. (1997)], [Mita et al.(1998)], [Mita et al.(1999)] [Xin et al.(2000)]. We explored those ideas in the nonlinear case. Reference [NL01] shows how a nonstandard H-infinity problem that

results from a specification to suppress constant disturbances, can be reduced to a standard problem with a smaller state dimension. This is achieved by reducing the order of the state feedback Hamilton-Jacobi differential equation.

Reference [NL02] adopts a quite different approach to the same problem. It is assumed that a nonlinear controller has already been designed for a given nonlinear plant, but that it does not necessarily suppress a constant disturbance or deliver zero error in tracking a constant reference input. It is shown, using singular perturbation theory [Kokotovic et al.(1986)], how to modify the given controller, through the addition of an integrator, in order to secure the desired constant disturbance rejection, while retaining the essential qualitative features of the original controller. The paper gives an example, applying the theory to a (multiple input multiple output) helicopter control problem [Koo and Sastry(1998)] obtained from the University of California, Berkeley.

Reference [NL03] produces a new control method for underactuated nonlinear systems, for variable constraint control. The results in the paper are applied to the posture control of free flying robots.

4.3 Safe adaptive control and iterative feedback tuning

In the early part of the contract on safe adaptive control and iterative feedback tuning, we were concerned with isolating those issues of fundamental importance in adaptive control which are applicable to nonlinear as well as linear adaptive control. There was consequently less focus on the design of particular algorithms than on what an adaptive control algorithm should do, and determining which difficulties also apply within existing linear theory. Since those difficulties represent considerable barriers in use, we concentrated on first repairing those difficulties before extending the theory to the nonlinear case.

Fundamental difficulties were identified in [AC01]. Since subsequent research was based on the issues raised there, we shall explain them in some detail, emphasising that safe adaptive control algorithms are needed in order to overcome those difficulties.

4.3.1 Problems of inexact modelling

It is common for an identification of the plant to be undertaken with a particular controller in the loop, either explicitly or implicitly, both in adaptive control algorithms, as well as in iterative identification and control design. The plant identification is always approximate rather than perfect. The quality of the approximation may be evaluated in terms of similarity between the closed-loop behaviour of the actual plant and current controller and that of the identified model and the current controller. If the behaviours are not very similar then the plant should be re-identified.

Assume that we have a good approximate plant model, but that the closed loop performance is poor. A traditional adaptive control algorithm would usually redesign the controller such that the new closed-loop comprising the identified model and the new controller has better performance, by either implicitly or explicitly using the identified model. The new controller is then implemented. Such a scheme relies on the implicit assumption that if the plant model yields a good approximation of closed loop behaviour with the original controller, then it will also do so for the new adjusted controller. However, such an assumption is not always valid, unless the change in the controller is small [Vinnicombe(1999)]. On the other hand, if the controller change is large, then even a model that results in a good approximation of closed-loop behaviour with the original controller may result in a very poor approximation with the new controller. In fact, the loop comprising the actual plant

and the new controller may be unstable, even if the loop comprising model and new controller has attractive performance.

It follows that adaptive control algorithms need to guard against the possibility that any controller changes invalidate the model that was used for its design. Although we have not described measures that quantify controller change, in the above description, we will return to this point in Section 4.3.4.

4.3.2 Transient instability

Many theorems in adaptive control texts [Goodwin and Sin(1984)], [Mareels and Polderman(1996)] assert that given certain assumptions, a given adaptive control algorithm will have the property that all signals in the closed-loop will remain bounded, and that convergence occurs as time tends to infinity. While superficially attractive, such theorems fail to address the quality of transient performance. In fact, it is possible that a controller is temporarily connected during the course the algorithm which, if left in place with unchanged parameters, would give an unstable closed-loop. In such a situation, the adaptive algorithm will detect such an instability and make corrective change to the controller. However, in the meanwhile, signals can become quite large. Such an adaptive control algorithm is fundamentally unsafe. In contrast, a safe adaptive control algorithm is one that does not result in such "transient instability".

4.3.3 Unattainable objectives

It is usual that closed-loop performance is part of the specification of an adaptive control problem, and that the plant is unknown to some degree. It is also well known that certain performance specifications are practically unobtainable for certain plants, even if obtainable in theory. For example, an open loop bandwidth of 1 Hz cannot be extended to a 1 kHz closed loop bandwidth in practice, even with the aid of feedback.

There is a risk in an adaptive control problem that not only is the performance objective unobtainable, but the initial uncertainty of the plant model means that this fact is unknown. An algorithm that does not detect that a particular controlled objective is impractical is likely to result in quite unacceptable performance. Safe adaptive algorithms need to indicate whether a specified closed loop performance objective is practically unobtainable. Very few adaptive algorithms do this.

We reiterate that the three problems referred to above all arise irrespective of the linearity of the system being controlled, so that understanding how to resolve these problems in the linear case is of value. For an introduction to these ideas see [AC01], [AC02].

4.3.4 A Windsurfer Approach to Safe Adaptive Control

Recent work by the chief investigator under the rubric "a windsurfer approach to adaptive control" addressed these issues implicitly, and demonstrated the safe control of plants with unmodeled resonances and of unknown model order [Lee et al.(1995)]. A more modern treatment is found in [AC03], which puts windsurfer ideas in the context of the above issues. An extension of [Lee et al.(1995)] to cope with open-loop unstable plants appears in [AC04].

In order to address the difficulties quantitatively, a particular metric, the *Nu-gap* (Vinnicombe) metric [Vinnicombe(1993)], [Vinnicombe(1999)] was used. During the contract we investigated the extension of the Vinnicombe metric to nonlinear operators with a view to extending the ideas of [Lee et al.(1995)] to nonlinear adaptive control in a quantitative as well as a qualitative fashion. Paper [AC05] sought to define a nonlinear generalization of the *Nu-gap* metric, but was partly incomplete as it relied on particular unproven conjectures. Paper [AC06] is complete, presenting an extension of the Vinnicombe metric to a pseudo-metric on Lipschitz continuous nonlinear operators (that is, those for which small changes in the input produce a small change in the output on a finite time interval). Although numerical calculations involving the nonlinear Vinnicombe metric will be difficult in general, there will almost certainly be classes of systems for which bounds are easily constructed, for example, those containing a simple nonlinearity, memoryless and sector-bounded. This is promising for a quantitatively-based approach to safe nonlinear adaptive control.

Papers [AC07]-[AC10] all deal with linear systems, and describe how one can do safe iterative modelling and control design. Problems are considered in which a largely unknown plant is given, together with a stabilizing controller. A new approximate model of the plant is constructed based on noisy, closed-loop measurements and a new controller is designed. This new controller is assured to be safe, in that attachment to the still partially unknown plant will not produce an unstable closed loop. A sequence of iterative identification and controller redesign ultimately leads to a satisfactory closed loop. In the event that the closed-loop specifications are too demanding, and cannot be achieved for the actual unknown plant, this will be indicated during the course of the algorithm, even were this fact initially unknown. In such a case, the algorithm will indicate that there is no value in further identification and controller redesign.

A nonlinear version of some of these ideas can be found in [AC12]. However the absence of excellent quantitative tools such as a nonlinear *Nu-gap* metric forms a roadblock to the practical application of this nonlinear theory.

The material on multiple model adaptive control in the section following draws heavily on a number of these ideas.

4.4 Multiple Model Adaptive Control

After approximately one year's work on the contract, Dr Moshfegh indicated an interest in multiple model adaptive control (see [Morse (1996)], [Morse (1998a)], [Morse (1998b)] and [Narendra and Balakrishnan (1997)]). There are several reasons why multiple model adaptive control should be contemplated.

1. Rather less is known about the performance of multiple model adaptive control algorithms than those based on continuously varying model parameters. This suggests a potentially rich source of new concepts and insights.
2. Many linear adaptive control algorithms rely on the fact that the model parameters appear linearly in the system equations. This is not a property of many nonlinear systems. Other approaches to nonlinear adaptive control need to be considered. Multiple model adaptive control is one such method that does not require that parameters appear linearly in particular equations.
3. Many adaptive control algorithms contain an explicit or implicit identification component. Identification of a continuously valued (but stationary) parameter usually results in a

parameter error variance that converges at a rate that is inversely proportional to the elapsed identification time. On the other hand, multiple model adaptive control scheme require distinguishing between multiple hypotheses. In that case, error rates (in the case of only two hypotheses, measured by false alarm or miss probabilities) decay exponentially with elapsed identification time. Thus, a model parameter value estimate that is in the vicinity of a correct value will be achieved faster by using multiple model adaptive control than by using a model with continuously variable parameters. Note that a multiple model adaptive control scheme can be augmented to allow parameter estimates to be tuned in a continuously variable manner, around the nominal value of a particular model from the multiple model set.

Dr Moshfegh highlighted that fact there did not appear to be a sound methodology for specifying the multiple models, given a parametric region defining the unknown model class. It is preferable to use the minimum number of models possible. Paper [MM01] describes how to construct a set of multiple models in an economic and systematic way by exploiting properties of the v -gap metric. The method produces a model set for an earlier published example both very quickly and in a systematic fashion, and that was much the same as when the multiple model set had been determined after much trial and error. The companion paper [MM02] compares two different switching logics for multiple model adaptive controllers, and focuses on the behaviour of the resulting closed-loop hybrid system. Methods of reliably detecting which model in the multiple model set is the most likely are the subject of [MM03].

Paper [MM04], to appear also, with modifications, as [MM05] combines the concept of multiple model adaptive control with safe switching. In a multiple model adaptive controller, at particular instants of time, one controller is replaced by another controller. Safe switching is switching which ensures that such replacements never produce a system which is even frozen closed-loop unstable.

Besides demonstrating the feasibility of such an algorithm, the paper shows that convergence is potentially slower. This is because more data is collected in order to ensure that a potential switch is safely justified than in the case when safety is disregarded. This paper also illustrates that without a safety constraint, transient instability can easily be encountered.

4.5 Hidden Markov Models

Hidden Markov models [Elliot et al. (1994)] are important, if for no other reason than that both many hybrid systems and multiple model adaptive control involve discrete states. Hidden Markov models (HMMs) of interest are those in which the state assumes one of a finite (countable) number of values. The value of the state evolves in a Markov fashion. Noisy measurements are available, although not necessarily of the state itself. Thus a measurement at one instant in time is insufficient to determine the state. The usual sorts of questions of filtering, prediction and smoothing arise.

Reference [HMM01] is a survey paper which connects hidden Markov models to Wiener and Kalman filtering. In particular, this survey exposes the important fact that a hidden Markov model filter will normally have an exponential rate of convergence. Therefore, no matter how it is initialized, the initial condition will be forgotten exponentially fast. In addition, round-off errors will not accumulate in a disastrous fashion, and outliers will eventually be forgotten. These are essential properties that a practical filter must have.

Although as in the case of Kalman and Wiener filters there is a simple formula for the filter equations, there is no simple formula for the performance of a hidden Markov model filter. There is no analogue of an error covariance that is available either a priori or a posteriori.

Work of others in the late 1990s established explicit performance formulae for Hidden Markov models in which the states changed very slowly [Golubev and Khasminskii(1998)] [Khasminskii and Zeitouni(1996)]. Nonrigorously, this corresponds to having a high signal to noise ratio at very low frequencies. Papers [HMM02] and [HMM03] obtain analogous results for smoothing as opposed to filtering, with a similar restriction on the models. These show that particularly in the case of a very slow state variation, there can be a substantial performance improvement achieved by using a fixed lag smoother as opposed to a filter, which may be traded off against the potential disadvantage that the estimates are not available as quickly.

It is widely recognized that communication channel limitations mean that analogue information cannot be sent over a channel in many practical circumstances. Instead discretization must occur before transmission. References [HMM04] and [HMM05] study the estimation of the state in a two-state hidden Markov model for two cases— a continuously distributed output, and a discretized version of that same output. The work addresses the two questions of how to choose the quantization levels in order to minimize the filtering error, and how the filtering error varies with noise for a different number of quantization levels. The papers only give partial answers to both those questions. At the moment, there is no straightforward or simply expressed rule of thumb that answers those questions.

Reference [HMM06] solves a very long standing problem in the area of hidden Markov models that is analogous to the realization problem in linear system theory. In the linear system realization problem, the Markov parameters of the system are given, and a state variable realization of the system is to be constructed. In the HMM realization problem, the probabilities of all finite length output strings are given, and a finite-state Markov process and a state-to-output mapping is to be constructed that generates an output process with the specified statistics. While not directly addressing either hybrid systems or multiple model adaptive control, it is regarded by some as a major advance in the general theory of hidden Markov models.

5 High Level Behaviour

This section outlines our work for the second research topic, namely the development of hierarchical control strategies. In our report [PR01] we identified some available theory [Wong and Wonham (1996)] for hierarchical supervisory control of *discrete event systems* (DES). While it is realistic to consider discrete dynamics on the high level as far as the control of several vehicles at the one time is concerned – low level dynamics in relevant detail are imposed by physical systems and are continuous by nature. This pointed our attention to the area of *hybrid systems*, that is systems which are composed from both continuous and discrete components, and the bulk of our work on high level problems has focused on hybrid systems.

In order to focus efforts, and in the light of the Navy requirements, we are posing a **Navy challenge problem** involving hierarchical and decentralized control, see also [PR02], Section 4.1. Consider three helicopters flying in line abreast. How can their configuration be changed so that they are flying in line, one behind another? Note that this is not a problem of maintaining a particular configuration, by making minor corrections in each vehicle. It involves a gross manoeuvre, during

which relative positions remain critical. Of course, one can overlay requirements such as the duration of the manoeuvres, and so on.

From a control engineering perspective, we suggest the design of individual low level controllers capable of executing basic manoeuvres on the individual helicopters, such as straight and level flight, constant rate climb, or constant rate turn. To achieve a desired change of formation, a high level controller is required to coordinate the manoeuvres. Our larger goal is the development of controller synthesis methods for the design and integration of both low level and high level controllers in such a way that the closed-loop behaviour is guaranteed to satisfy its performance specifications.

Obviously, such methods must take into account the dynamics of each of the components involved as well as the interaction between them. Thus any suitable overall model is expected to be of high complexity – giving rise to an implementation challenge. A key issue will be to consider the hierarchical structure, thus giving rise a control theoretic challenge. As we have identified before, classical control has very little to say about this type of problem, despite the fact that the control of each individual helicopter necessarily involves classical control concepts. Instead, methods of hybrid systems need to be brought to bear.

The choice of a rather specific problem serves two purposes. First, it clearly emphasizes the relevance of our work to high level problems of interest to the Navy. Second, it provides guidance through scientific roadblocks, as unsolved subproblems are brought to light. Outcomes are reported under the following headings, each motivated by our Navy challenge problem:

- A modal logic framework for hybrid systems.
- Robust control of hybrid systems.
- Modular control of hybrid systems.
- Hierarchical control of hybrid systems.

5.1 A modal logic framework for hybrid systems

Hybrid systems are heterogeneous dynamical systems characterized by interacting continuous and discrete dynamics, and typically arise in the embedded software control of physical processes. Such mathematical models have proved fruitful in a great diversity of engineering applications, including automated transportation, robotics, and automated manufacturing. In particular, the above Navy challenge problem sets up a hybrid system and that motivates our interest in hybrid systems in the context of our ONR contract.

In this section we report our work on a quite general *modal logic* based framework for the synthesis of hybrid systems. The general idea to apply *formal methods* to dynamical systems was originally developed for the analysis of computer hardware and software systems which can be modelled as purely discrete finite state machines, but some ideas have subsequently been extended and adapted to deal with hybrid systems. One challenge here is to make key properties of continuous dynamics accessible to a formal framework. The reference [Davoren and Gore (2001)] provides an axiomatisation of semi-continuity properties, a key tool for reasoning formally about continuous or hybrid dynamics. The dominant trend in formal methods for hybrid systems is to use the framework of *temporal logic* where the emphasis is on *formal verification*; that is giving a formal proof that a

system fulfils a specification [Alur et al. (1996), Alur et al. (2000), Manna and Pnueli (1993a), Manna (1998)]. In our work we combine ideas from both control theory and computer science to develop a *modal logic* based approach for the *formal synthesis* of hybrid systems. This concept was suggested by in [HL01] and since then has been considerably extended. One major outcome is a synthesis algorithm that solves a general class of hybrid control problems [HL02]. The algorithm is stated within our formal framework and exploits the power of modal logic.

The hybrid control problem under consideration is stated from a traditional control theory point of view: given a switched continuous plant, construct a switching controller so that the resulting closed-loop system is guaranteed to satisfy a list of performance specifications. Here, the plant consists of a finite number of continuous systems $x' = F_c(x)$ over a common state space X , a subset of n -dimensional Euclidean Space, indexed by symbols c elements of C in a finite (discrete) control alphabet. The controller exhibits discrete dynamics, realized on a finite state space Q , and includes an output mapping from Q to the control alphabet C . The controller must decide when to switch its discrete state q to another state p , and output a new control symbol based on its continuous measurement of the plant state X . This decision mechanism is represented by a controller transition relation α a subset of $Q \times X \times Q$, which determines two sorts of regions of the plant state space: regions in which the controller grants permission to *stay* in a discrete state, and regions in which the controller grants permission to *switch* from one discrete state to another. The closed-loop dynamics can be represented by the widely accepted hybrid automaton model, where the so called *mode invariants* and *guard regions* correspond to staying regions and switching regions, respectively.

The types of qualitative behavioural specifications we address go beyond the class of *safety*, *invariance* and *reachability* properties, which are the sole or primary focus of much of the current work on hybrid controller synthesis [Asarin et al. (2000a), Lygeros et al. (1999), Tomlin et al. (2000)]. Safety properties are usually formulated as negative reachability assertions, of the form: no hybrid trajectory starting in a given set of initial states will ever enter a set *Bad*, where *Bad* is a proscribed set of plant states. In our target class of control problems, we additionally address positive or active behavioural requirements. We deal with a very general class of *event sequence* properties, of the form: all hybrid trajectories must traverse in a prescribed order through the blocks of a given finite partition of the plant state space. This gives a general-purpose way of specifying the attainment of local goals along the course of hybrid trajectories, and integrating the type of event sequence specifications examined in DES approaches to hybrid systems [Koutsoukos et al. (2000), Horn and Ramadge (1995), Moor and Raisch (1999a)]. We also address the two basic forms of *liveness* properties: that all hybrid trajectories can be extended indefinitely, to make infinitely many discrete changes of state, and that all hybrid trajectories be *non-Zeno* (so not make infinitely many discrete switches in finite real time).

Our essential idea is that in designing and constructing the switching controller for a given plant and given specifications, one needs to reason about *sets of plant states*, and build up more complicated sets of states by applying various operators arising from the flows and the specification data. Following [HL01], we use *modal logic* as a clean and elegant formalism in which to conduct such reasoning about sets of states, and to custom-design operators on sets tailored to the specifications. The logic gives us the technical tools with which to formulate a general and finitely terminating synthesis algorithm which applies uniformly to arbitrary differential equations $x' = F_c(x)$, subject only to standard assumptions on the existence and uniqueness of solutions, with finite termination analytically derived from an assumption of compactness. By formulating these constructions of complex sets of states in the language of modal logic, we gain the immediate pay-off that the *correctness* of the synthesis procedure – that any controller generated by the procedure does indeed ensure that the closed-loop system satisfies all the performance specifications —can transparently

transparently be shown to be a *formal deductive consequence* of a theory of modal formulas that are true of that hybrid automaton model purely in virtue of the construction.

The modal logic framework also gives us a clean way to separate out the determination of what sets need to be computed, and the structure and correctness of the abstract solution algorithm, from the distinct issue of how and when such an algorithm can be *effectively implemented*. Effective implementation requires a finitary symbolic means of representing set of states A , a subset of n -dimensional Euclidean space, with respect to which the Boolean and modal logic operators can be effectively evaluated, and furthermore, the representation of sets must be *decidable* in the sense that it can be determined by finite computation whether distinct representations are semantically equal. These are the fundamental issues for the application and development of *symbolic model checking tools* for hybrid and real-time systems [HL01], see also [Alur et al. (2000), Asarin et al. (2000a), Henzinger and Majumdar (2000)]. There are two main approaches to effective implementation, based on *exact symbolic representations* of state sets A or on *approximated representation* of sets of states, working with under- or over-approximations. Recent contributions to approximation methods for the basic forwards and/or backwards reachability operators of differential equations (and differential inclusions) are variously based on boxes [Asarin et al. (2000a) Bournez et al. (1999), Maler and Dang (1998), Moor and Raisch (1999a)], polyhedra [Chutinan and Krogh (1998)] or ellipsoids [Kurzanski and Variaya (2000)]. The publication [HL03] continues research in box shaped approximations of various reachability operators. Each of these approximation techniques apply to arbitrary linear differential equations, and in principle, any of them could serve as a basis for approximated versions of the modal operators used in our abstract synthesis algorithm. In [HL04], we discuss fundamental properties of reachability operators and their approximations in the presence of uncertainty.

5.2 Robust control of hybrid systems

The principal motivation for robust control designs immediately applies to hybrid control systems: we ask for a controller that enforces a desired performance specification in the presence of plant uncertainty. In our Navy challenge problem, continuous low-level controllers implement elementary manoeuvres, and we may ask for a robust design that addresses for example a range of weather conditions, a range of battle damage conditions or a range of load conditions. While such a robust low-level controller will maintain performance up to a certain degree, it cannot be expected that the continuous closed-loop is completely independent on weather conditions. Here, a sensible goal for high-level controller synthesis is the ability to handle the remaining parameter uncertainty in the continuous closed-loop system. In our careful discussion of a typical example we document that without any further precautions a hybrid control design can fail to exhibit even elementary robustness properties.

We formally address the problem of robust hybrid controller design within our modal logic framework, as outlined in Section 5.1. Our crucial observation is that various classes of uncertainty—including the traditional plant parameter uncertainty—can be expressed in terms of *metric tolerance relations*, and the effect of these relations on sets of states can be captured by modal logic operators. By using these notions of metric tolerance, we are able to cleanly formulate and prove several forms of robustness or tolerance properties for our synthesis algorithm. Our result is that not only is it the case that all hybrid trajectories of the nominal closed-loop system H meet the given specifications, but in addition, all hybrid trajectories arising from certain *bounded variations* of H will still meet those specifications. The variation classes we consider arise by allowing a bounded degree of tolerance of sensor and actuator imprecision, [HL02] as well as bounded variations in the differential equations defining the plant, where the variation depends continuously on a parameter.

Both the latter, traditional formulation of robustness in terms of plant variation, as well as our notions of sensor and actuator tolerance, fall within a framework of robustness concepts for hybrid automata proposed by Horn and Ramadge in [Horn and Ramadge (1995)].

5.3 Modular control of hybrid systems

The construction of an overall supervisory controller by combining a number of individual supervisors is referred to as *modular supervisory control*. In the situation of our Navy challenge problem, we consider two high-level control problems separately: (i) the design of a controller C_1 that avoids collisions of the three helicopters; (ii) the design of a controller C_2 that adjusts the relative positions of individual helicopters according to the desired change of formation. Assuming that both control problems have been solved, the question arises whether it is possible to combine C_1 and C_2 such that both specifications are enforced simultaneously; that is whether we can compose a controller that achieves the desired change of formation while in the same time collisions are avoided. In our particular example, the synthesis of C_2 can be further decomposed into smaller subproblems by considering each one of the helicopters. Obviously, such decomposition is not possible in the construction of the collision avoidance controller. Thus, the synthesis of C_1 is expected to be computationally expensive. On the other hand, the problem of collision avoidance is of a general interest which is not restricted to the particular formations under consideration. The concept of modularity will enable us to recycle C_1 for various versions of C_2 which address various formation reconfigurations.

From a more general point of view, we ask for *sufficient conditions* that allow two supervisors, each enforcing a particular specification, to be combined to enforce both specifications simultaneously. The motivation for attempting modular control is twofold: (i) the synthesis of individual supervisors and their subsequent combination might be computationally less expensive than the direct synthesis of an overall controller; (ii) based on the concept of modular control, one may set up a "library" of supervisors, each geared towards a specific task for a given plant; depending on the particular application situation, the appropriate controllers can then be simply retrieved from the library and run in parallel to solve the problem at hand. In the field of DES theory, modularity has been studied (for example) [Wonham (1999), Ramadge and Wonham (1989), Rudie and Wonham (1992)] and our strategy is to extend these results to general classes of hybrid control systems. We use the framework set up in earlier work [Moor and Raisch (1999b)], where we discuss the problem of supervisor synthesis for hybrid systems with discrete external signals. This work is set within Willems' behavioural systems theory, and extends the core of Ramadge and Wonham's DES theory to the considered class of hybrid systems. Our recent results in [HL06] show how the concept of modularity as it is stated in can be applied to hybrid systems with discrete external signals.

5.4 Hierarchical control of high-order hybrid systems

A scenario that has been commonly used as a motivation for hybrid control consists of a continuous plant model, a finite number of continuous controllers and a discrete supervisor which acts on quantized measurement information (events) by switching between the continuous controllers. It is clear that this scenario exhibits a (two-level) hierarchical structure: the continuous feedback loops can be interpreted as lower-level control, the supervisor to be designed as a higher-level controller. In the scope of our Navy challenge problem, the plant corresponds to the helicopters, the continuous controllers implement the elementary manoeuvres, and the discrete supervisor is supposed to

coordinate these manoeuvres according to the desired change of formation. In [HL08] we exploit this two-level hierarchical architecture in the course of supervisory controller synthesis.

Given a continuous plant with discrete external signals we ask for a high-level supervisory controller that enforces a language inclusion specification. This problem has been considered extensively and solutions are typically based on computing a suitable (that is, conservative) *discrete abstraction* for the continuous part of the overall system; for example [Koutsoukos et al. (2000), Moor and Raisch (1999b)] The crucial computational challenge in this step is to reliably estimate sets of continuous states reachable under continuous flows from different sets of initial conditions. For fairly large classes of continuous dynamics this can be done by employing a regular quantisation grid in the continuous state space; for example [Bournez et al. (1999), Franke et al. (2000), Lygeros et al. (1999)]. (Technically, [Lygeros et al. (1999)] restates the reachability problem as a partial reachability problem as a partial differential equation, which is then to be solved numerically.)

Clearly, this puts a rather stringent limit on the problem state dimension. In the context of control, hierarchies are mostly introduced to “break” a complex problem into a number of more tractable problems [Caines and Wei (1998), Farzoli et al. (1999), Pappas et al. (2000), Raisch et al. (2000), Raisch et al. (2001), Wong and Wonham (1996)] and hence to reduce the overall ‘solution effort’. We therefore expect that the hierarchical structure in our set-up can be exploited to significantly reduce the computational burden in the abstraction step. More precisely, we argue that the presence of low-level controllers may considerably reduce the dimension of the part of the continuous state space that is relevant for the abstraction step.

Using a grid partitioning of the n -dimensional state space, the number of discrete states in the abstraction depends exponentially on n . We identify a general class of low-level control goals that is characterised by an m -dimensional stable component of the state variable. This enables us to effectively reduce the dimension of the state space to $n - m$. Computational advantage is then gained for two reasons: first, the lower dimensional grid consists of significantly fewer cells; second, the long term continuous dynamics can be approximated by a reduced model. This second aspect requires a detailed analysis of the continuous feedback loops, and we give such an analysis for the situation of linear time invariant differential equations in [HL08]. Our method is reliable in the sense that it is still guaranteed that the original system will only evolve on trajectories that are generated by the abstraction. This condition is crucial when employing the discrete abstraction as a basis for supervisory controller synthesis. As a benchmark, we applied our method to the design of a start-up procedure of a distillation-column [H107].

6 Conclusion

Individuals working on the contract

Principal Investigator: B. D. O. Anderson (1998-2001)

Associate Investigators

T. S. Brinsmead (1999-2001), J. Davoren (1999-2001), F. De Bruyne (1998-1999), S. Dey (1998-1999), L. Shue (2000), T. Moor (2000-2001), K. Wong (1999)

Students

N. Linard (1998-1999), S. Su (1999-2001), M. Smith (2000-2001), L. Shue (1998-2000)

References

Papers Produced under Contract

- [AC01] B. D. O. Anderson and M. Gevers. Fundamental problems in adaptive control. In D. Normand-Cyrot, editor, *Perspectives in Control, Theory and Applications*. Springer-Verlag, 1998a.
- [AC02] B. D. O. Anderson and M. Gevers. Fundamental problems of control of adaptive processes. In Proc. 37th Conference of the Japanese Society of Instrumentation and Control Engineers. Chiba, Japan, 1998b.
- [AC03] B. D. O. Anderson. Windsurfing approach to iterative control design. In P. Albertos, editor, *Iterative Identification and Control Design*. Springer, 2002 to appear.
- [AC04] W. S. Lee, I. M. Y. Mareels, and B. D. O. Anderson. Iterative identification and two step control design for partially unknown unstable plants. *International Journal of Control*, 740 (1): 0 43--57, 2001.
- [AC05] B. D. O. Anderson and F. De Bruyne. On a nonlinear generalisation of the v gap metric. Proc. 38th IEEE Conf. on Decision and Control, pages 3851--3856, 1999.
- [AC06] B. D. O. Anderson, T. S. Brinsmead, and F. De Bruyne. The Vinnicombe metric for nonlinear operators. *IEEE Transactions on Automatic Control*, 2002. to appear.
- [AC07] B. D. O. Anderson, X. Bombois, M. Gevers, and C. Kulscar. Caution in iterative modeling and control design. *Proceedings IFAC Workshop on Adaptive Systems in Control and Signal Processing*, pages 13--19, 1998.
- [AC08] X. Bombois, B. D. O. Anderson, and M. Gevers. Mapping parametric confidence ellipsoids to Nyquist plane for linearly parametrized transfer functions. In G. C. Goodwin, editor, *Model Identification and Adaptive Control.*, pages 53--71. Springer-Verlag, London, 2000a.
- [AC09] M. Gevers, B. D. O. Anderson, and B. Codrons. Issues in modelling for control. *Proceedings American Control Conference*, pages 1615--1619, 1998.
- [AC10] R. R. Bitmead, B. D. O. Anderson, M. Gevers, and L. C. Kammer. Cautious controller tuning. Technical report, RSISE, Australian National University, Acton ACT 0200, Australia, 2000. submitted to Automatica.
- [AC11] X. Bombois, M. Gevers, G. Scorletti, and B. D. O. Anderson. Controller validation for stability and performance based on an uncertainty region designed from an identified model. *Proceedings IFAC Workshop on System Identification, 2000b*.
- [AC12] F. De Bruyne, B. D. O. Anderson, and I. D. Landau. Recursive iterative feedback tuning. *IEEE Conf. on Decision and Control*, pages 3804--3809, 1999a.
- [HL01] J. Davoren and A. Nerode. Logics for hybrid systems. *Proceedings of the IEEE*, 88:0 985--1010, 2000.

- [HL02] J. Davoren and T. Moor. Logic-based design and synthesis of controllers for hybrid systems. Technical report, RSISE, Australian National University, 2000. Submitted for publication.
- [HL03] T. Moor and J. Raisch. Approximation of multiple switched flow systems for the purpose of control synthesis. In Proc. of the 39th International Conference on Decision and Control, CDC'00. IEEE Press, 2000.
- [HL04] M. A. Smith, T. Moor, and B. D. O. Anderson. Restricted reachability operators: safe approximations in the presence of disturbances. Technical report, RSISE, Australian National University, 2001. Submitted to the 15th IFAC World Congress.
- [HL05] T. Moor and J. Davoren. Robust controller synthesis for hybrid systems using modal logic. In M. D. Benedetto and A. Sangiovanni-Vincentelli, editors, *Hybrid Systems: Computation and Control* (HSCC'01), volume 2034 of LNCS, pages 433--446. Springer-Verlag, 2001.
- [HL06] T. Moor, J. Davoren, and J. Raisch. Modular supervisory control of a class of hybrid systems in a behavioural framework. In *Proceedings of the European Control Conference 2001*, pages 870--875, 2001a.
- [HL07] T. Moor, J. Raisch, and A. Itigin. Discrete approximation and control of high-order nonlinear continuous systems. Technical report, RSISE, Australian National University, 2001c. Submitted to the 15th IFAC World Congress.
- [HL08] T. Moor, J. Raisch, and J. Davoren. Computational advantages of a two-level hybrid control architecture. Technical report, RSISE, Australian National University, 2001b. Accepted for presentation at the 40th International Conference on Decision and Control, CDC'01.
- [HL09] B. D. O. Anderson, F. De Bruyne, S. Dey, and K. Wong. Ensuring robustness in hybrid control systems. Technical report, RSISE, Australian National University, 1999b. Submitted.
- [HMM01] B. D. O. Anderson. From Wiener to hidden Markov models: Common threads of filtering and smoothing. *IEEE Control Systems Magazines*, pages 41-51, 1999a.
- [HMM02] L. Shue, B. D. O. Anderson, and F. De Bruyne. Asymptotic smoothing errors for hidden markov models. *IEEE Trans. Signal Processing*, 280 (12):0 3289--3302, 2000.
- [HMM03] S. Dey. Reduced-complexity filtering for partially observed nearly completely decomposable markov chains. *Proc. Math. Theory of Networks and Systems*, 2000.
- [HMM04] L. Shue, S. Dey, B. D. O. Anderson, and F. De Bruyne. Remarks on error due to quantisation of a 2-state hidden markov models. *Proc. IEEE Conf. on Decision and Control*, pages 4123--4128, 1999.
- [HMM05] L. Shue, B. D. O. Anderson, and F. De Bruyne. On state-estimation of a 2-state hidden markov model with quantisation. *IEEE Trans. Signal Processing*, 49 (1): 202ff, 2001.
- [HMM06] B. D. O. Anderson. The realization problem for hidden Markov models. *Mathematics of Control, Signals and Systems*, 12:80-120, 1999b.

- [ID01] B. D. O. Anderson. From Youla-Kucera to identification, adaptive and robust control. *Automatica*, 34(12):1485-1505, 1998.
- [ID02] N. Linard, B. Anderson, and F. De Bruyne. Coprimeness properties of nonlinear fractional system realizations. *Systems and Control Letters*, 34:0 265--271, 1998.
- [ID03] N. Linard, B. Anderson, and F. De Bruyne. Identification of a nonlinear plant under nonlinear feedback using left-coprime fractional based representations. *Automatica*, 35:655-667, 1999.
- [ID04] K. Fujimoto, B. D. O. Anderson, and F. De Bruyne. Differentially coprime kernel representations and closed-loop identification of nonlinear systems. *Proceedings of the American Control Conference*, 2000a.
- [ID05] K. Fujimoto, B. D. O. Anderson, and F. De Bruyne. A paremetrization for closed-loop identification of nonlinear systems based on differentially coprime kernel representations. *Proceedings of the American Control Conference*, 2000.
- [ID06] K. Fujimoto, B. D. O. Anderson, and F. De Bruyne. A paremetrization for closed-loop identification of nonlinear systems based on differentially coprime kernel representations. Technical report, RSISE, Australian National University, submitted to *Automatica*, 2001.
- [ID07] F. De Bruyne, B. D. O. Anderson, M. Gevers, and N. Linard. On a closed-loop identification with a tailor-made parametrization. *Proceedings of the American Control Conference*, pages 3177-3182, 1998.
- [ID08] F. De Bruyne, B. D. O. Anderson, M. Gevers, and N. Linard. Gradient expressions for a closed-loop identification scheme with a tailor-made parametrization. *Automatica*, 35(11):1867-1871, 1999.
- [ID09] I. D. Landau, B. D. O. Anderson, and F. De Bruyne. Closed-loop output error identification algorithms for nonlinear plants. *Proc. 38th IEEE Conf. on Decision and Control*, pages 606-611, 1999.
- [ID10] F. De Bruyne, I. D. Landau, and B. D. O. Anderson. Closed-loop output error identification algorithms of nonlinear plants using kernel representations. *Proc. 39th IEEE Conf. on Decision and Control*, 2000.
- [ID11] I. D. Landau, B. D. O. Anderson, and F. De Bruyne. Recursive identification algorithms for continuous time nonlinear plants operating in closed-loop. *Automatica*, 36:0 469--475, 2001.
- [ID12] I. D. Landau, B. D. O. Anderson, and F. De Bruyne. Algorithms for identification of continuous time nonlinear systems: a passivity approach. part I: Identification in open-loop operation. part ii: Identification in closed-loop operation. In A. Isidori, F. Lamnabhi-Lagarigue, and W. Respondek, editors, *Nonlinear Control in the Year 2000*, pages 13--44. Springer-Verlag, 2000.
- [ID13] F. De Bruyne, B. D. O. Anderson, and N. Linard. Implementation issues for a nonlinear version of the Hansen scheme. In *Information, Decision and Control*, pages 295--300, Adelaide, Australia, 1999b.

[ID14] B. Codrons, B. D. O. Anderson, and M. Gevers. Closed-loop identification with an unstable or nonminimum phase controller. *Proceedings IFAC Symposium on System Identification*, 2000.

[ID15] F. De Bruyne, B. D. O. Anderson, and N. Linard. The Hansen scheme revisited. In *Proc. IEEE Conf. on Decision and Control*, pages 706-711, Tampa, Florida, USA, 1998.

[MM01] B. D. O. Anderson, T. S. Brinsmead, F. De Bruyne, J. Hespanha, D. Liberzon, and A. S. Morse. Multiple model adaptive control, part 1: Finite controller coverings. *International Journal of Robust and Nonlinear Control*, 10(11-12):909-927, 2000b.

[MM02] J. P. Hespanha, D. Liberzon, A. S. Morse, B. D. O. Anderson, T. S. Brinsmead, and F. De Bruyne. Multiple model adaptive control, part 2: Switching. *International Journal on Robust and Nonlinear Control*, 11(5): 479-496, 2001.

[MM03] F. De Bruyne and S. Dey. Change detection and its application to adaptive control. *Proceedings of the American Control Conference* 2000.

[MM04] B. D. O. Anderson, T. S. Brinsmead, J. Hespanha, D. Liberzon, and A. S. Morse. Safe switching in multi-controller implementation. *IFAC Symposium on Adaption and Learning in Control and Signal Processing*, 2001a.

[MM05] B. D. O. Anderson, T. S. Brinsmead, J. Hespanha, D. Liberzon, and A. S. Morse. Multiple model adaptive control with safe switching. *International Journal of Adaptive Control and Signal Processing*, 2001b. to appear.

[NL01] S. W. Su, B. D. O. Anderson, and T. S. Brinsmead. Robust disturbance suppression for nonlinear systems based on H-infinity design. *Proc. IEEE Conf. on Decision and Control*, pages 3013--3018, 2000.

[NL02] S. W. Su, B. D. O. Anderson, and T. S. Brinsmead. Constant disturbance suppression for nonlinear systems design using singular perturbation theory. *Proc. 40th IEEE Conf. on Decision and Control*, 2001. submitted.

[NL03] T. Ikeda, T.-K. Nam, T. Mita, X. Xin, and B. D. O. Anderson. Variable constraint control of nonholomic free flying robots—mechanical design and convergence. *Proc. 38th IEEE Conf. on Decision and Control*, 1999.

[PR01] B. D. O. Anderson, F. De Bruyne, S. Dey, and K. Wong. Interim progress report: Development of design methodologies for robust, adaptive, autonomous, hierarchical systems. Technical report, RSISE, Australian National University, Acton ACT 0200, Australia, April 1999b.

[PR02] B. D. O. Anderson, T. S. Brinsmead, J. Davoren, T. Moor, and S. Su. Interim progress report: Development of design methodologies for robust, adaptive, autonomous, hierarchical systems. Technical report, RSISE, Australian National University, Acton ACT 0200, Australia, March 2000a.

Other References

- [Alur et al.(1996)] R. Alur, T. Henzinger, and P.-H. Ho. Automatic symbolic verification of embedded systems. *IEEE Transactions on Software Engineering*, 22: 181– 201, 1996.
- [Alur et al. (2000)] R. Alur, T. Henzinger, G. Lafferriere, and G. Pappas. Discrete abstractions of hybrid systems. *Proceedings of the IEEE*, 88, 2000.
- [Anderson and Moore (1979)] B. D. O. Anderson and J. B. Moore. *Optimal Filtering*. Prentice-Hall International, Englewood Cliffs, NJ, c1979.
- [Asarin et al. (2000a)] E. Asarin, O. Bournez, T. Dang, O. Maler, , and A. Pnueli. Effective synthesis of switching controllers for linear systems. *Proceedings of the IEEE*, 88:0 1011–1025, 2000.
- [Bournez et al. (1999)] O. Bournez, O. Maler, and A. Pnueli. Orthogonal polyhedra: representation and computation. In F. Vaandrager and J. van Schuppen, editors, *Hybrid Systems: Computation and Control* (HSCC'99), LNCS 1569, pages 46--60. Springer-Verlag, 1999.
- [Caines and Wei (1998)] P. Caines and Y. Wei. Hierarchical hybrid control systems: a lattice theoretic formulation. *IEEE Transactions on Automatic Control*, 43:4:501–508, 1998.
- [Chutinan and Krogh (1998)] A. Chutinan and B. Krogh. Computing polyhedral approximations to flow pipes for dynamic systems. In *Proc. of the 37th International Conference on Decision and Control, CDC'98*. IEEE Press, 1998.
- [Dasgugta and Anderson (1996)] S. Dasgugta and B. Anderson. A parametrization for the closed-loop identification of nonlinear time-varying systems. *Automatica*, 1996.
- [Davoren and Gore (2001)] J. Davoren and R. Gore. Bimodal logics for reasoning about continuous dynamics. Technical report, RSISE, Australian National University, 2001. Accepted for publication in *Advances in Modal Logic*.
- [Elliot et al. (1994)] R. J. Elliot, L. Aggoun, and J. B. Moore. *Hidden Markov Models*. Applications of Mathematics. Springer, NY, 1994.
- [Farzzoli et al. (1999)] E. Farzzoli, M. Dahleh, and E. Feron. A hybrid control architecture for aggressive maneuvering of autonomous helicopters. In *Proceedings of the 38th International Conference on Decision and Control, CDC'99*, pages 2471--2476. IEEE Press, 1999.
- [Forsell and Ljung (1999)] U. Forssell and L. Ljung. Closed-loop identification revisited. *Automatica*, 350 (7):0 1215--1241, 1999.
- [Franke et al. (2000)] D. Franke, T. Moor, and J. Raisch. Discrete supervisory control of switched linear systems. *at-Automatisierungstechnik*, 48:9:0 461–467, 2000.
- [Franklin et al. (1994)] G. F. Franklin, J. D. Powell, and A. Emami-Naeni. *Feedback Control of Dynamic Systems*. Addison-Wesley, third edition edition, 1994.

[Gevers et al. (1999)] M. Gevers, B. Codrons, and F. De Bruyne. Model validation in closed-loop. In *Proc. American Control Conf.*, pages 326--330, San Diego, USA, 1999.

[Golubev and Sin Khasminskii (1998)]. G. Golubev and R. Khasminskii. Asymptotic optimal filtering for a hidden Markov model. *Math. Methods Statist*, 70 (2):192--209, 1998.

[Goodwin and Sin (1984)] G.C. Goodwin and K. S. Sin. *Adaptive Filtering Prediction and Control*. Prentice-Hall, 1984.

[Henzinger and Majumdar (2000)] T. Henzinger and R. Majumdar. A classification of symbolic transition systems. In *Proc. of 17th International Symposium on Theoretical Aspects of Computer Science (STACS'00)*, LNCS. Springer-Verlag, 2000.

[Hansen et al. (1989)] F. R. Hansen, G. F. Franklin, and R. L. Kosut. Closed-loop identification via the fractional representation: Experiment design. *Proc. American Control Conference*, pages 386--391, 1989.

[Horn and Ramadge (1995)] C. Horn and P. Ramadge. Robustness issues for hybrid systems. In *Proceedings of the 34th International Conference on Decision and Control, CDC'95*, pages 1467--1472. IEEE Press, 1995.

[Jiang and Praly (1998)] Z.-P. Jiang and L. Praly. Design of robust adaptive controllers for nonlinear systems with dynamic uncertainties. *Automatica*, 1998.

[Kokotovic et al. (1986)] P. Kokotovic, H.K. Khalil, and J.O'Reilly. *Singular Perturbation Methods in Control: Analysis and Design*. Academic Press Inc, U.S.A., 1986.

[Koo and Sastry (1998)] T.J. Koo and S. Sastry. Output tracking control design of a helicopter model based on approximate linearization. *Proceedings of the 37th Conference on Decision and Control*, 1998.

[Koutsoukos et al. (2000)] X. Koutsoukos, P. Antsaklis, J. Stiver, and M. Lemmon. Supervisory control of hybrid systems. *Proceedings of the IEEE*, 88:0 1026--1049, 2000.

[Khasminskii and Zeitouni (1996)] R. Khasminskii and O. Zeitouni. Asymptotic filtering for finite state Markov chains. *Stoch. Proc. Their Appl.*, 630 (1):0 1--10, 1996.

[Krstic et al. (1995)] M. Krstic, I. Kanellakopoulos, and P. Kokotovic. *Nonlinear and Adaptive Control Design*. John Wiley & Sons, NY, 1995.

[Kuo (1991)] B. Kuo. *Automatic Control*. Prentice Hall, 1991.

[Kurzhanski and Varaiya (2000)] A. Kurzhanski and P. Varaiya. Ellipsoidal techniques for reachability analysis. In B. Krogh and N. Lynch, editors, *Hybrid Systems: Computation and Control (HSCC'00)*, LNCS 1790. Springer-Verlag, 2000.

[Lee et al. (1995)] W. S. Lee, B. D. O. Anderson, I. M. Y. Mareels, and R. L. Kosut. On some key issues in the windsurfer approach to adaptive robust control. *Automatica*, 31(11): 1619-1636, 1995.

[Ljung(1987)] L. Ljung. *System Identification, A Theory for the User*. Prentice-Hall International, NJ, 1987.

[Ljung and Forssell (1999)] L. Ljung and U. Forssell. A alternative motivation for the indirect approach to closed loop identification. *IEEE Trans. Automatic Control*, AC-440 (11): 2206—2209, 1999.

[Lygeros et al. (1999)] J. Lygeros, C. Tomlin, and S. Sastry. Controllers for reachability specifications for hybrid systems. *Automatica*, 35:0 349--370, 1999.

[Maler and Dang (1998)] O. Maler and T. Dang. Reachability analysis via face lifting. In T. Henzinger and S. Sastry, editors, *Hybrid Systems: Computation and Control (HSCC'98)*, LNCS 1386, pages 96--109. Springer-Verlag, 1998.

[Manna (1998)] H. S. Z. Manna. Deductive verification of hybrid systems using STeP. In T. Henzinger and S. Sastry, editors, *Hybrid Systems: Computation and Control (HSCC'98)*, LNCS, pages 305--318. Springer-Verlag, 1998.

[Manna and Pnueli (1993a)] Z. Manna and A. Pnueli. Models for reactivity. *Acta Informatica*, 30: 609--678, 1993.

[Mareels and Polderman (1996)] I. M. Y. Mareels and J. W. Polderman. *Adaptive systems : an introduction*. Birkhauser, Boston, 1996.

[Mita et al. (1998)] T. Mita, M. Hirata, K. Murata, and H. Zhang. H-infinity control versus disturbance-observer based control. *IEEE Trans. on Industrial Electronics*, pages 488--495, 1998.

[Mita et al. (1997)] T. Mita, X. Xin, and B. D. O. Anderson. Extended H-infinity control— solving H-infinity servo and estimation problems. *Proc. 36th IEEE Conf. on Decision and Control*, pages 4653—4658, 1997.

[Mita et al. (1999)] T. Mita, X. Xin, and B. D. O. Anderson. Extended H-infinity control & H-infinity with unstable weights. *Automatica*, 1999. to appear.

[Moore and Raisch (1999a)] T. Moor and J. Raisch. Discrete control of switched linear systems. In *Proceedings of the European Control Conference 1999*, 1999a.

[Moor and Raisch (1999b)] T. Moor and J. Raisch. Supervisory control of hybrid systems within a behavioural framework. *Systems and Control Letters*, 38: 157—166, 1999b.

[Morari and Zafiriou (1989)] M. Morari and E. Zafiriou. *Robust Process Control*. Prentice Hall, NJ, 1989.

[Morse (1996)] A. S. Morse. Supervisory control of families of linear set-point controllers—part 1: Exact matching. *IEEE Trans. Auto. Control*, 41 (10): 1413--1431, 1996.

[Morse (1998a)] A. S. Morse. A bound for the disturbance-to-tracking-error gain of a supervised set-point control system. In D. Normand-Cyrot, editor, *Perspectives in Control, Theory and Applications*, pages 23--41. Springer-Verlag, 1998.

[Morse (1998b)] A. S. Morse. Control using logic-based switching. In A. Isidori, editor, *Trends in Control: A European Perspective*, pages 69—113. Springer-Verlag, 1998b.

[Narendra and Balakrishnan (1997)] K. S. Narendra and J. Balakrishnan. Adaptive control using multiple models. *IEEE Trans. Auto. Control*, 42(2):171—187, February 1997.

[PAPPAS ET AL. (2000)] G. Pappas, G. Lafferriere, and S. Sastry. Hierarchically consistent control systems. *IEEE Transactions on Automatic Control*, 45:6: 1144—1160, 2000.

[Raisch et al. (2000)] J. Raisch, A. Itigin, and T. Moor. Hierarchical control of hybrid systems. In S. Engell, S. Kowalewski, and J. Zaytoon, editors, *Proc. 4th International Conference on Automation of Mixed Processes: Dynamic Hybrid Systems*, pages 67–72. Shaker Verlag, 2000.

[Raisch et al. (2001)] J. Raisch, A. I. J.M., and T. Moor. Hierarchical strategies for hybrid process control problems. In *Proceedings of the European Control Conference 2001*, pages 2534–2539, 2001.

[Ramadge and Wonham (1989)] P. Ramadge and W. Wonham. Modular control of discrete event systems. *Maths. of Control, Signals Systems*, 1:1:13–30, 1989.

[Rudie and Wonham (1992)] K. Rudie and W. Wonham. Think globally, act locally: decentralized supervisory control. *IEEE Transactions on Automatic Control*, 37:11:1692—1708, 1992.

[Sepulchre et al. (1996)] R. Sepulchre, M. Jankovic, and P. Kokotovic. *Constructive Nonlinear Control*. Springer Verlag, New York, 1996.

[Sjöberg et al. (1995)] J. Sjöberg, Q. Zhang, L. Ljung, A. Benveniste, B. Delyon, P.-Y. Glorennec, H. Hjalmarsson, and A. Juditsky. Nonlinear black-box modeling in system identification: a unified overview. *Automatica*, 31 (12): 1691–1724, 1995.

[Tomlin et al. (2000)] C. Tomlin, J. Lygeros, and S. Sastry. A game theoretic approach to controller design for hybrid systems. *Proceedings of the IEEE*, 88, 2000. To appear; Special Issue on Hybrid Systems.

[Van den Hof and Schrama (1995)] P. van den Hof and R. Schrama. Identification and control, closed-loop issues. *Automatica*, 1995.

[Van den Hof and Schrama (1995)] A. J. Van der Schaft. *L_2 Gain and Passivity Techniques in Nonlinear Control*. Springer Verlag, New York, 1996.

[Vidyasagar (1985)] M. Vidyasagar. *Control System Synthesis: A Factorization Approach*. MIT Press, Cambridge, Massachusetts, 1985.

[Vinnicombe (1993)] G. Vinnicombe. Frequency domain uncertainty and the graph topology. *IEEE Trans. Auto. Control*, 38(9): 1371–1383, 1993.

[Vinnicombe (1999)] G. Vinnicombe. *Uncertainty and Feedback: H-infinity Uncertainty and the Nu-gap metric*. Imperial College Press, London, 2001.

[Wong and Wonham (1996)] K. Wong and W. Wonham. Hierarchical control of discrete event systems. *Discrete Event Dynamic Systems*, 6: 241–273, 1996.

[Wonham (1999)] W. Wonham. Notes on control of discrete event systems. Technical report, Department of Electrical & Computer Engineering, University of Toronto, 1999. Downloadable on <http://odin.control.toronto.edu/DES/>.

[Xin et al. (2000)] X. Xin, B. D. O. Anderson, and T. Mita. Complete solution of the 4-block H_∞ control problem with infinite and finite $j\omega$ -axis zeros. *Int. J. of Robust and Nonlinear Control*, 100 (2):0 59--81, 2000.

[Youla et al. (1976a)] D. C. Youla, J. J. Bongiorno, Jr., and H. A. Jabr. Modern Wiener-Hopf design of optimal controllers, part I: the single-input-output case. *IEEE Trans. Auto. Control*, AC-21: 3–13, 1976a.

[Youla et al. (1976b)] D. C. Youla, H. A. Jabr, and J. J. Bongiorno, Jr. Modern Wiener-Hopf design of optimal controllers, part II: the multivariable case. *IEEE Trans. Auto. Control*, AC-21: 319–338, 1976b.

[Zhou et al. (1996)] K. Zhou, J. C. Doyle, and K. Glover. *Robust and Optimal Control*. Prentice Hall, NJ, 1996.